UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
15/098,910	04/14/2016	David A. Gilbert	1145-012/20161032US	9819

137131 7590 07/31/2020

Red Hat, Inc and Withrow & Terranova, PLLC
106 Pinedale Springs Way
Cary, NC 27511

EXAMINER	
PAPERNO, NICHOLAS A	

ART UNIT	PAPER NUMBER
2132	

NOTIFICATION DATE	DELIVERY MODE
07/31/2020	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

patents@wt-ip.com

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

Ex parte DAVID A. GILBERT

Appeal 2019-002763
Application 15/098,910
Technology Center 2100

Before ST. JOHN COURTENAY III, LARRY J. HUME, and PHILLIP A. BENNETT, *Administrative Patent Judges*.

COURTENAY, *Administrative Patent Judge*.

DECISION ON APPEAL

Appellant[1] appeals under 35 U.S.C. § 134 from a Final Rejection of claims 1–21. We have jurisdiction over the pending claims under 35 U.S.C. § 6(b).

We affirm in part.

[1] We use the word "Appellant" to refer to "applicant" as defined in 37 C.F.R. § 1.42(a) (2018). According to Appellant, the real party in interest is Red Hat, Inc. *See* Appeal Br. 1.

STATEMENT OF THE CASE[2]

Introduction

Appellant's claimed invention relates generally to "fault tolerant computing devices and, in particular, to reducing memory inconsistencies between two synchronized computing devices." Spec. 1.

Illustrative Independent Claim 1

1. A first computing device comprising:

a random access memory (RAM) comprising a plurality of memory pages;

a processor device coupled to the RAM; and

a first hypervisor module that interfaces with the processor device and is to iteratively:

determine that content of a memory page of the plurality of memory pages has been modified;

send the content of the memory page to a second hypervisor module on a second computing device;

identify, for verification, at least one other memory page of the plurality of memory pages that was previously sent to the second computing device;

generate a verification value based on content of the at least one other memory page; and

[L] send the verification value and a memory page identifier that identifies the at least one other memory page to the second hypervisor module on the second computing device *without sending the content of the at least one other*

[2] We herein refer to the Final Office Action, mailed May 3, 2018 ("Final Act."); Appeal Brief, filed Nov. 5, 2018 ("Appeal Br."); the Examiner's Answer, mailed Dec. 20, 2018 ("Ans."); and the Reply Brief, filed Feb. 20, 2019 ("Reply Br.").

> *memory page to allow the second hypervisor module to verify that a copy of the at least one other memory page maintained by the second hypervisor module is identical to the at least one other memory page*.

Appeal Br. 25, "CLAIMS APPENDIX" (bracketed lettering added and dispositive disputed limitation "**L**" emphasized).

Evidence

The prior art relied upon by the Examiner as evidence is:

Name	Reference	Date
Lu et al.	US 7,440,982 B2	Oct. 21, 2008
Venkitachalam et al.	US 2010/0318991 A1	Dec. 16, 2010
Mangtani et al.	US 2013/0232498 A1	Sept. 5, 2013
Bissett et al.	US 8,812,907 B1	Aug. 19, 2014
Bhargava et al.	US 2015/0143064 A1	May 21, 2015
Hunter, "vSphere Replication Synchronization Types," June 19, 2015, https://blogs.vmware.com/vsphere/2015/06/vsphere-replication-synchronization-types.html		

Rejections

Rej.	Claims Rejected	35 U.S.C. §	Reference(s)/Basis
A	1, 3,[3] 5–7, 9–12, 14–17[4]	103	Venkitachalam et al. ("Venkitachalam"), Hunter, Bissett et al. ("Bissett"), Lu et al. ("Lu")
B	2, 4, 13	103	Venkitachalam, Hunter, Bissett, Lu, Bhargava et al. ("Bhargava")
C	8	103	Venkitachalam, Hunter, Bissett, Lu, Mangtani et al. ("Mangtani")
D	18, 19	103	Venkitachalam, Bissett, Bhargava
E	20	103	Venkitachalam, Bissett, Bhargava, Mangtani
F	21	103	Venkitachalam, Bissett, Bhargava, Lu

Analysis

Issue under § 103

Issue: Under 35 U.S.C. § 103, we focus our analysis on the following argued limitation regarding Rejection A of independent claims 1 and 12.

Did the Examiner err by finding that Venkitachalam, Hunter, Bissett, and Lu collectively teach or suggest negative limitation "**L**":

> [**L**] send the verification value and a memory page identifier . . . *without sending the content of the at least one other memory page to allow the second hypervisor module to verify that a*

[3] Claim 3 was omitted from the heading for Rejection A (Final Act. 2), but a detailed statement of rejection under Rejection A is included on page 7 of the Final Action. Appropriate correction has been made above.

[4] Claims 14 and 16 were omitted from the heading for Rejection A (Final Act. 2), but a detailed statement of rejection under Rejection A is included on page 11 of the Final Action. Appropriate correction has been made above.

> *copy of the at least one other memory page maintained by the second hypervisor module is identical to the at least one other memory page*[,]

within the meaning of independent claim 1?[5]

The Examiner reads the disputed dispositive negative limitation **L** of claim 1 on Lu, at Fig. 2, column 6, lines 10–63. *See* Final Act. 6.

The Examiner finds, "[s]ince the metadata has to be obtained by checking the indexes, it can be retrieved (meaning it would have to be sent to the component doing the comparison) without having to also obtain the data sets that generated." Final Act. 6.

Appellant disagrees with the Examiner's underlying factual findings and ultimate legal conclusion of obviousness. On page 15 of the Appeal Brief, Appellant traverses the teachings of Lu, and contends:

> Again, Lu discloses a process for "verifying stored data" (which, by definition, is not a volatile memory page) (Lu, Abstract). Thus, Appellant submits that "stored data" has nothing to do with, and fails to teach or suggest "*a random access memory (RAM) comprising a plurality of memory pages.*" Lu discloses a process wherein a first set of metadata is generated based on a first copy of data, and the metadata is "stored to an index" (Lu, col. 6, lines 13-15, 32-33). A second set of metadata is generated based on a second copy of data and is "stored to an index" (Lu, col. 6, lines 37-40, 50-51). The first set of metadata and the second set of metadata are compared to determine if the second set of metadata is an accurate copy of the first set of metadata (Lu, col. 6, lines 52-67). Thus, Lu discloses a process for

[5] We give the contested claim limitations the broadest reasonable interpretation (BRI) consistent with the Specification. *See In re Morris*, 127 F.3d 1048, 1054 (Fed. Cir. 1997).

> determining whether two copies of data are identical. Lu contains no disclosure about hypervisor modules or memory pages. Thus, Lu cannot disclose" . . . *without sending the content of the* ***at least one other memory page*** *to allow the* ***second hypervisor module*** *to verify that a copy of the* ***at least one other memory page*** *maintained by the* ***second hypervisor module*** *is identical to* ***the at least one other memory page."***

Appeal Br. 15.

The Examiner disagrees with Appellant, and further explains the basis for the rejection in the Answer:

> Lu is relied on to teach two things: a memory that is RAM and the sending of a verification value *without having to send the contents being verified.* The memory that the data is stored on does not affect the verification process. Just because data in the references is stored on disk or tape as stated in the appellant's arguments does not mean the processes would be different if they are stored on RAM. The data would still need to be identified as being modified or not, a verification value would still need to be generated and sent to the appropriate place to be compared with another value to determine if a synchronization operation needs to occur. The purpose of Lu is to ensure that the copy of the first set of data was done without any error (Abstract). This means that the method of Lu is comparing verification values to check to see if different sets of stored data are identical. Also the fact that Lu does not teach virtual machines or hypervisors does not mean that it is irrelevant and cannot be combined with the other references. As stated previously, Venkitachalam is relied upon to teach the majority of the architecture and system that is implementing the synchronization procedure. Venkitachalam does not specify the memory type which is why Lu is used because it mentions a specific memory that is a RAM.

Ans. 11–12 (emphasis added).

As an initial matter of claim construction, we focus our analysis on the negative claim limitation: "*without sending the content of the at least one other memory page to allow the second hypervisor module to verify that a copy of the at least one other memory page maintained by the second hypervisor module is identical to the at least one other memory page*." Claim 1 (emphasis added).

Appellant indicates that paragraph 34 of the Specification provides written description support for the negative limitation. *See* Appeal Br. 2 (under "Summary of the Claimed Subject Matter"). Paragraph 34 describes, in pertinent part: "Note that the memory pages 24-2, 24-7 themselves may not be sent." Spec. 34.

Based upon our review of the record, we find the Examiner has not shown that the disputed negative limitation of claim 1 is taught or suggested by Lu individually, or in combination with the other cited references. We find the closest teaching is Lu's description at column 6, lines 44–48: "In some embodiments, the second set of metadata generated in a partial copy operation, wherein, a set of data is read, metadata based on the set of data is generated, but no copy of the data file is made."

However, for essentially the same reasons argued by Appellant (Appeal Br. 15), we find this description neither teaches nor suggests the negative limitation of sending a verification value and a memory page identifier "*without sending the content of the at least one other memory page to allow the second hypervisor module to verify that a copy of the at least one other memory page maintained by the second hypervisor module is identical to the at least one other memory page*." Claim 1 (emphasis added).

As noted above, the Examiner only relies upon the Lu reference, to teach or suggest the recited negative limitation of claim 1. *See* Final Act. 6–7. Therefore, we are constrained on this record to reverse the Examiner's obviousness Rejection A of independent claim 1.[6] Independent method claim 12, also rejected under Rejection A, recites the disputed negative limitation **L** of claim 1 using identical language.

Accordingly, for the same reasons, we also reverse the Examiner's Rejection A of independent claim 12. Because we have reversed the Examiner's Rejection A of independent claims 1 and 12, we also reverse the Examiner's Rejection A of each associated dependent claim.

---

[6] In the event of further prosecution, *including any review prior to allowance*, we leave it to the Examiner to consider a rejection under 35 U.S.C. § 112, first paragraph (written description). In particular, we note the negative limitation added by amendment during prosecution: "without sending the content of the at least one other memory page . . . ." Claim 1. We do not find a *reason to exclude* sending the content of the at least one other memory page in the originally filed Specification. Although the amended claim itself provides a reason to exclude ("*to allow the second hypervisor module to verify that a copy of the at least one other memory page maintained by the second hypervisor module is identical to the at least one other memory page*," this explanation was not provided in the original claims. Our reviewing court provides applicable guidance: "Negative claim limitations are adequately supported when the specification describes a reason to exclude the relevant limitation. Such written description support need not rise to the level of disclaimer. In fact, it is possible for the patentee to support both the inclusion and exclusion of the same material." *Santarus, Inc. v. Par Pharmaceutical, Inc*., 694 F.3d 1344, 1351 (Fed. Cir. 2012). *See also* MPEP § 2173.05(i) ("Any negative limitation or exclusionary proviso must have basis in the original disclosure . . . . The mere absence of a positive recitation is not basis for an exclusion."). Although the Board is authorized to reject claims under 37 C.F.R. § 41.50(b), no inference should be drawn when the Board elects not to do so. *See* Manual of Patent Examining Procedure (MPEP) § 1213.02.

Rejection B of Claims 2, 4, and 13, and
Rejection C of Claim 8

In light of our reversal of Rejection A of independent claims 1, and 12, *supra*, we also reverse obviousness rejection B of dependent claims 2, 4, and 13, and obviousness Rejection C of dependent claim 8.

On this record, the Examiner has not shown how the additionally-cited secondary references overcome the aforementioned deficiencies with respect to the base combination of Venkitachalam, Hunter, Bissett, and Lu, as discussed above regarding independent claims 1 and 12, as rejected under Rejection A.

Accordingly, we are constrained on this record to reverse the Examiner's obviousness rejections A, B, and C of claims 1–17 on appeal.

Rejection D of Claims 18 and 19

Regarding independent claim 18, Appellant contend:

> The Patent Office relies upon Bissett for disclosing the determination of a set of memory pages that has been modified, and for the identification of a second set of memory pages, and refers to its rejection of claim 6. However, as Appellant noted above with regard to claim 6, the fact that Bissett maintains a list of memory pages that have been modified does not mean that Bissett actually discloses identifying a second group of **unmodified** memory pages in addition to the memory pages that have been modified. Perhaps the Patent Office is suggesting that it would be possible to do so; however, the possibility of doing so does not actually disclose doing so.

Appeal Br. 22.

The Examiner concludes that the identification of the first group of *modified* memory pages means that the second group of *unmodified* memory

pages is also *necessarily* identified, because what is not selected as "modified" is by default "unmodified." This is because any given memory page is either modified or not modified — a binary outcome. Ans. 22.

We note "inherency may supply a missing claim limitation in an obviousness analysis." *PAR Pharmaceutical, Inc. v TWI Pharmaceuticals, Inc*. 773 F.3d 1186, 1194-95 (Fed Cir. 2014).

We find a preponderance of the evidence supports the Examiner's legal conclusion of obviousness, because there are only two finite states of the memory pages, i.e., "modified" or "unmodified." Our reviewing court provides further applicable guidance: When "the problem is known, the possible approaches to solving the problem are known and finite, and the solution is predictable through use of a known option," a solution that is *obvious to try* may indeed be obvious. *Abbott Labs. v. Sandoz, Inc.*, 544 F.3d 1341, 1351 (Fed. Cir. 2008) (emphasis added), citing *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 421 (2007).

Appellant further contends:

> Moreover, the existence of two different memory stores on two different virtual machines fails to teach or suggest determining both a set of memory pages that has been modified and the identification of a set of other memory pages on the same computing device. The Patent Office relies on column 7, lines 3-17 of Bissett as disclosing the generation of a single verification value based on a plurality of memory pages. Appellant disagrees, and submits that column 7, lines 3-17 of Bissett discloses the generation of a checksum for disk data. Appellant submits that the generation of a checksum for disk data fails to teach or suggest the generation of a single verification value based on a plurality of memory pages.

Appeal Br. 22.

The Examiner explains:

> The reason that the data on the other disk can be considered the different group is because the mirror copy is not updated with the same frequency, hence the need to keep track of dirty and clean pages. Given that the data can be of some size, which can be multiple pages, the identification of the data on the mirrored drive that is then used to create a checksum (validation value) that is used for comparison does read on the cited limitations.

Ans. 22–23.

We are not persuaded by Appellant's arguments and agree with the Examiner's findings, because we find Bissett's "checksum" teaches, or at least suggests, the "verification value" recited in independent claim 18. *See* Bissett, col. 4, ll. 52–55, col. 7, ll. 5–11. We further find Appellant is arguing the references separately, as the Examiner cites Venkitachalam, not Bissett, for teaching memory pages. *See* Final Act. 2–3.

Combinability under 35 U.S.C. § 103

As pertinent to claims 18 and 19, the Appellant also challenges the Examiner's motivation to combine the references. *See* Appeal Br. 22, last paragraph.

It is our view that the Examiner's proffered combination of the respective cited features of Venkitachalam, Bissett, and Bhargava would have merely realized a predictable result.[7] *See* Final Act. 15–17.

[7] In *KSR*, the Court stated: "[t]he combination of familiar elements according to known methods is likely to be obvious when it does no more than yield predictable results." *KSR*, 550 U.S. at 416.

Thus, it would have been obvious to one of ordinary skill in the art to **modify** the pre-existing checksum and identification of page groups to obtain the predictable results of **determining** that the content of each memory page of a first set of memory pages of the plurality of memory pages has been modified; **send** the content of each memory page in the first set of memory pages to a second hypervisor module on a second computing device; randomly **identify** a second set of memory pages from the plurality of memory pages; **generate** a verification value based on the content of each memory page in the second set of memory pages; and **send** a group of memory page identifiers, each memory page identifier identifying a corresponding one of the memory pages in the second set of memory pages, and the verification value, to the second hypervisor module. *See supra* n.7. *See* Final Act. 16.

"[W]hen a patent 'simply arranges old elements with each performing the same function it had been known to perform' and yields no more than one would expect from such an arrangement, the combination is obvious." *KSR*, 550 U.S. at 417 (quoting *Sakraida v. Ag Pro, Inc*., 425 U.S. 273, 282 (1976)).

Further, in *KSR*, 550 U.S. at 418, the Court stated that when considering obviousness that "the analysis need not seek out precise teachings directed to the specific subject matter of the challenged claim, for a court can take account of the inferences and creative steps that a person of ordinary skill in the art would employ." "When there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to pursue the known options within his or her technical grasp." *See KSR*, 550 U.S. at 421.

Moreover, it is sufficient that references suggest doing what Appellant did, although the Appellant's particular purpose was different from that of the references. *See In re Heck*, 699 F.2d 1331, 1333 (Fed. Cir. 1983) (citing *In re Gershon*, 372 F.2d 535, 539 (CCPA 1967)). For a prima facie case of obviousness to be established, *the reference need not recognize the same problem solved by the Appellants. See In re Kemps*, 97 F.3d 1427, 1430 (Fed. Cir. 1996). [N]either the particular motivation nor the avowed purpose of the [Appellant] controls" when performing an obviousness analysis. *KSR*, 550 U.S. at 419.

This reasoning is applicable here. In reviewing the record, we are not persuaded the Examiner erred because Appellant does not point to any evidence of record that shows combining the teachings of Venkitachalam, Bissett, Bhargava in the manner proffered by the Examiner (Final Act. 15–17) would have been "uniquely challenging or difficult for one of ordinary skill in the art" or would have "represented an unobvious step over the prior art." *Leapfrog Enters., Inc. v. Fisher-Price, Inc*., 485 F.3d 1157, 1162 (Fed. Cir. 2007) (citing *KSR*, 550 U.S. at 418).

Appellant has not identified knowledge gleaned only from the present application that was not within the level of ordinary skill at the time the claimed invention was made. *See In re McLaughlin*, 443 F.2d 1392 (CCPA 1971). Moreover, Appellant has not provided any objective evidence of secondary considerations (e.g., unexpected results), which our reviewing court guides "operates as a beneficial check on hindsight." *Cheese Sys., Inc. v. Tetra Pak Cheese & Powder Sys., Inc.*, 725 F.3d 1341, 1352 (Fed. Cir. 2013).

We find one of ordinary skill in the art would have understood that randomly identifying a set of memory pages from the plurality of memory pages, and determining that the content has been modified, would have merely realized a predictable result. Thus, we find the Examiner sets forth a sufficient rational underpinning explaining why an artisan would have combined the teachings of Venkitachalam, Bissett, and Bhargava. *See* Final Act. 15–17.

Therefore, on this record, and based upon a preponderance of the evidence, we are not persuaded of error regarding the Examiner's underlying factual findings and ultimate legal conclusion of obviousness regarding Rejection D of independent claim 18.

Accordingly, we sustain the Examiner's Rejection D of independent claim 18. To the extent Appellant argues dependent claim 19 separately, we note our discussion above regarding the two states of a memory page: "modified" or "unmodified," as being *obvious to try* given the finite number of two possible states. Therefore, we also sustain the Examiner's Rejection D of dependent claim 19.

Rejection E of Claim 20 and Rejection F of claim 21

Appellant does not advance separate, substantive arguments for dependent claims 20 and 21, rejected under Rejections E and F, respectively. Arguments not made are waived. *See* 37 C.F.R. § 41.37(c)(1)(iv). Accordingly, we sustain the Examiner's obviousness Rejections E and F of claims 20 and 21, respectively.

CONCLUSIONS

The Examiner erred in rejecting claims 1–17, as being obvious under 35 U.S.C. § 103, over the cited combinations of references.

The Examiner did not err in rejecting claims 18–21, as being obvious under 35 U.S.C. § 103, over the cited combinations of references.

DECISION SUMMARY

In summary:

Claims Rejected	**35 U.S.C. §**	**Reference(s)/Basis**	**Affirmed**	**Reversed**
1, 3, 5–7, 9–12, 14–17	103	Venkitachalam, Hunter, Bissett, Lu		1, 3, 5–7, 9–12, 14–17
2, 4, 13	103	Venkitachalam, Hunter, Bissett, Lu, Bhargava		2, 4, 13
8	103	Venkitachalam, Hunter, Bissett, Lu, Mangtani		8
18, 19	103	Venkitachalam, Bissett, Bhargava	18, 19	
20	103	Venkitachalam, Bissett, Bhargava, Mangtani	20	
21	103	Venkitachalam, Bissett, Bhargava, Lu	21	
Overall Outcome			18–21	1–17

FINALITY AND RESPONSE

No time period for taking any subsequent action in connection with this appeal may be extended under 37 C.F.R. § 1.136(a)(1)(iv). *See* 37 C.F.R. § 41.50(f).

AFFIRM IN PART